Request For Comments: 787

A. Lyman Chapin
July 1981

Subject: Connectionless Data Transmission Survey/Tutorial

From: A. Lyman Chapin

The attached paper on connectionless data transmission is being distributed to the members of a number of US organizations that are involved or interested in the development of international data communication standards. Following a review period ending September 1, 1981, a revised version of the paper - incorporating comments and suggestions received from reviewers - will be considered by the American National Standards Institute (ANSI) committee responsible for Open Systems Interconnection (OSI) Reference Model issues (ANSC X3T5). If approved, it will then be presented to the relevant International Organization for Standardization (ISO) groups as the foundation of a US position recommending the incorporation of connectionless data transmission by the Reference Model and related OSI service and protocol standards.

Your comments on the paper, as well as an indication of the extent to which the concepts and services of connectionless data transmission are important to you and/or your organization, will help to ensure that the final version reflects a true US position. They should be directed to the author at the following address:

A. Lyman Chapin
Data General Corporation MS E111
4400 Computer Drive
Westborough, MA 01580

(617) 366-8911 x3056

X3S33/X3T56/81-85
X3T5/81-171
X3T51/81-44
X3S37/81-71R

WORKING PAPER
This document has not been reviewed or approved by the appropriate Technical Committee and does not at this time represent a USA consensus.

Connectionless Data Transmission

A. Lyman Chapin

22 May 1981 Revision 1.00

ABSTRACT

The increasingly familiar and ubiquitous Reference Model of Open Systems Interconnection, currently being considered by the International Organization for Standardization (ISO) for promotion to the status of a Draft International Standard, is based on the explicit assumption that a "connection" - an association between two or more communicating entities, possessing certain characteristics over and above those possessed by the entities themselves - is required for the transfer of data in an Open Systems Interconnection (OSI) environment. Although the connection-oriented model of communications behavior has proven to be an extremely powerful concept, and has been applied successfully to the design and implementation of protocols and systems covering a wide range of applications, a growing body of research and experience suggests that a complementary concept - connectionless data transmission - is an essential part of the Open Systems Interconnection architecture, and should be embraced as such by the OSI Reference Model. This paper explores the concept of connectionless data transmission and its relationship to the more familiar concepts of connection-oriented data transfer, developing a rationale for the inclusion of the connectionless concept in the Reference Model as an integral part of the standard description of the OSI architecture.

1 Introduction

Over the past three years, a number of national and international standards organizations have expended the time and efforts of a great many people to achieve a description of an architectural Reference Model for interconnecting computer systems considered to be "open" by virtue of their mutual use of standard communication protocols and formats. The current description, the Reference Model of Open Systems Interconnection (RM/OSI)[1], is generally accepted by the International Organization for Standardization (ISO), the International Telephone and Telegraph Consultatitive Committee (CCITT), the European Computer Manufacturer's Association (ECMA), and many national standards bodies, including the American National Standards Institute (ANSI), and has progressed to the status of a Draft Proposed Standard (DP7498) within ISO. It describes the concepts and principles of a communications architecture organized hierarchically, by function, into seven discrete layers, and prescribes the services that each layer must provide to the layer immediately above it (the uppermost layer provides its services to user applications, which are considered to be outside of the Open Systems Interconnection environment). Building on the services available to it from the next-lower layer, each layer makes use of standard OSI protocols which enable it to cooperate with other instances of the same layer (its "peers") in other systems (see Figure 1). This technique of grouping related functions into distinct layers, each of which implements a set of well-defined services that are used by the layer above, partitions a very complex, abstract problem - "how can the components of a distributed application, operating in potentially dissimilar environments, cooperate with each other?" - into a number of more manageable problems that enjoy a logical relationship to each other and can individually be more readily understood.

The Reference Model was developed to serve as a framework for the coordination of existing and future standards designed to facilitate the interconnection of data processing systems. The purpose of OSI is to enable an end-user application activity (called an "application process") located in a system that employs OSI procedures and protocols (an "open" system) to communicate with any other appication process located in any other open system. It is not the intent of OSI to specify either the functions or the implementation details of systems that provide the OSI capabilities. Communication is achieved by mutual adherence to agreed-upon (standardized) services and protocols; the only thing that an OSI entity in a given layer in one system needs to know about an OSI entity in the same layer

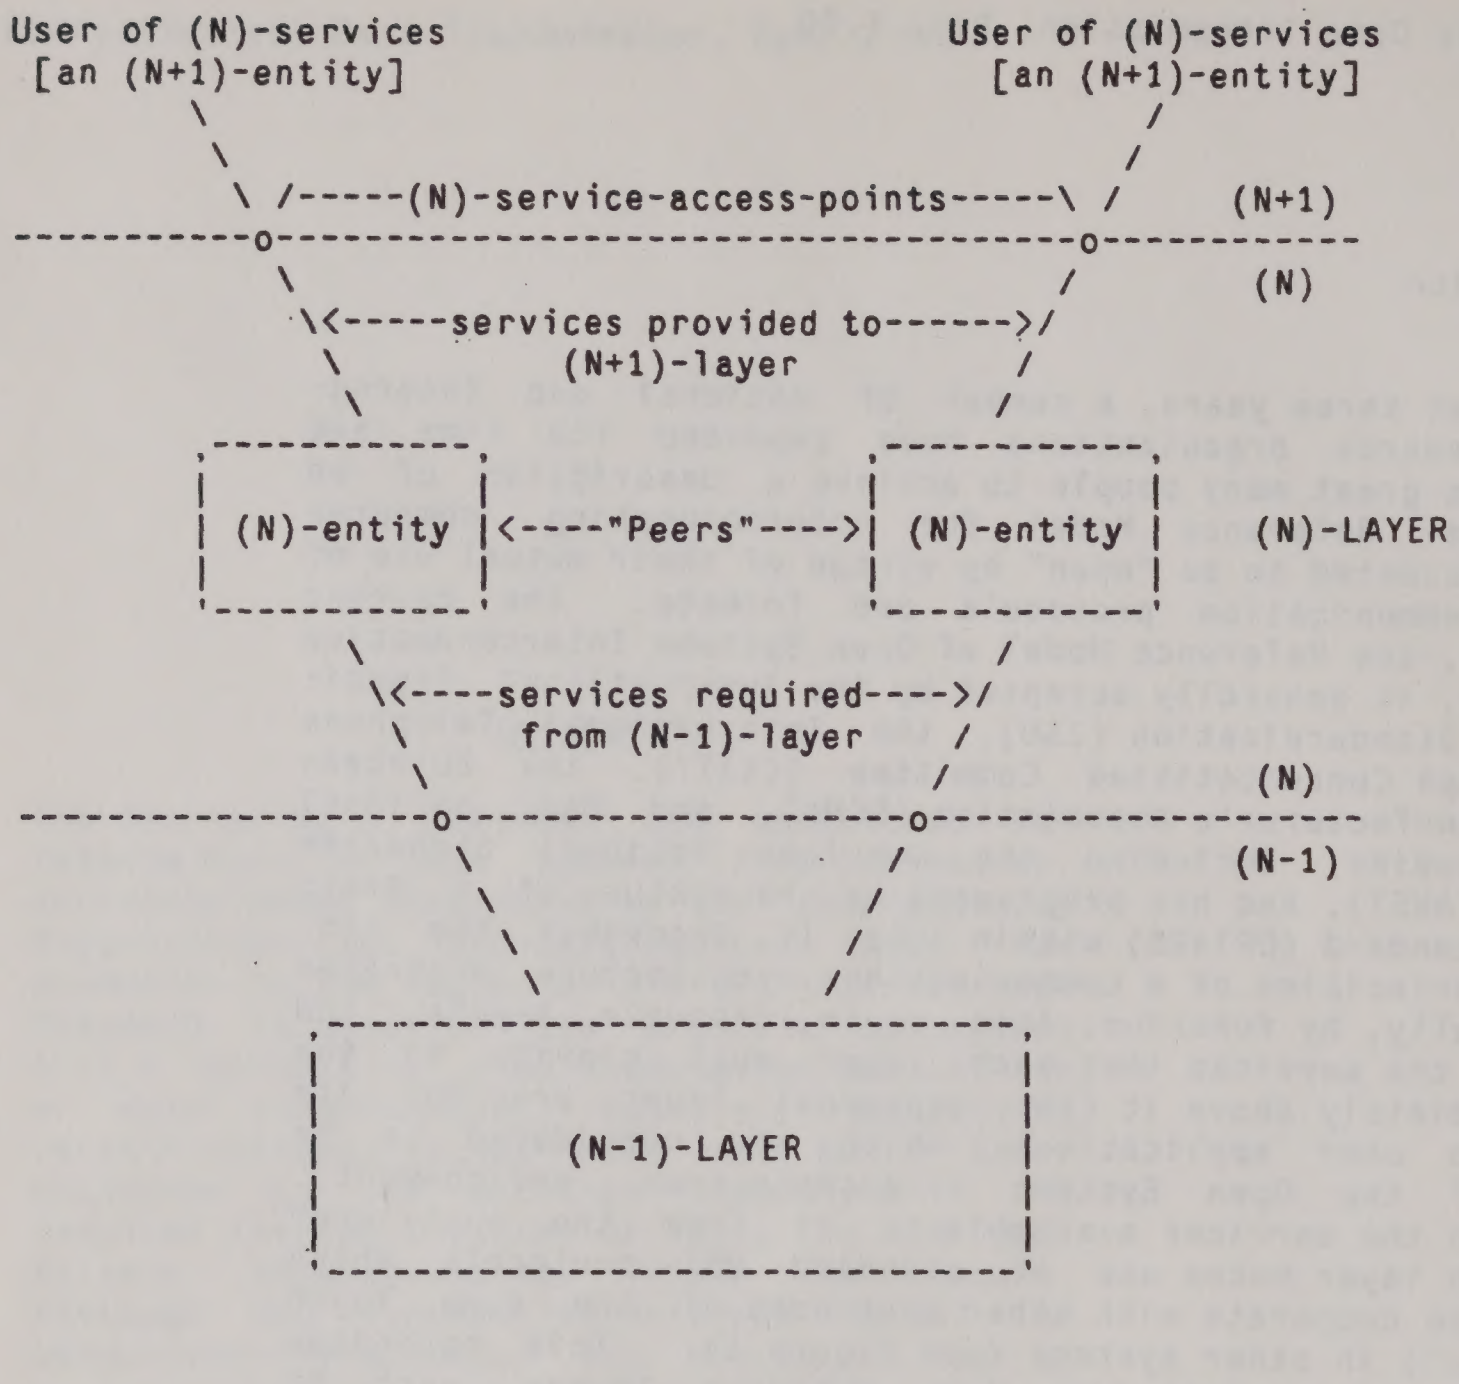

FIGURE 1 - General Model of an OSI Layer

A Note on OSI Terminology

The construction of a formal system, such as the architecture of Open Systems Interconnection, necessarily involves the introduction of unambiguous terminology (which also tends to be somewhat impenetrable at first glance). The terms found here and in the text are all defined in an Appendix. The "(N)-" notation is used to emphasize that the term refers to an OSI characteristic that applies to each layer individually. The "(N)-" prefix stands in generically for the name of a layer; thus, "(N)-address", for example, refers abstractly to the concept of an address associated with a specific layer, while "transport-address" refers to the same concept applied to the transport layer.

of another system is how the other entity behaves, not how it is implemented. In particular, OSI is not concerned with how the interfaces between adjacent layers are implemented in an open system; any interface mechanism is acceptable, as long as it supports access to the appropriate standard OSI services.

A major goal of the OSI standardization effort is generality. Ideally, the Reference Model should serve as the common architectural framework for many different types of distributed systems employing a wide range of telecommunication technologies, and certainly an important measure of the success of OSI will be its ability to apply the standard architecture across a broad spectrum of user applications. The way in which the Reference Model has developed over the past four years reflects an awareness of this goal (among others): the process began with the identification of the essential concepts of a layered architecture, including the general architectural elements of protocols, and proceeded carefully from these basic principles to a detailed description of each layer. The organization of the current Reference Model document [1] exhibits the same top-down progression. At the highest level, three elements are identified as basic to the architecture[1]:

a) the application processes which exist within the Open Systems Interconnection environment;

b) the connections which join the application processes and permit them to exchange information; and

c) systems.

The assumption that a connection is a fundamental prerequisite for communication in the OSI environment permeates the Reference Model, and is in fact one of the most useful and important unifying concepts of the architecture. A growing number of experts in the field, however, believe that this deeply-rooted connection orientation seriously and unnecessarily limits the power and scope of the Reference Model, since it excludes a large class of applications and implementation technologies that have an inherently connectionless nature. They argue that the architectural objectives of the Reference Model do not depend on the exclusive use of connections to characterize all OSI interactions, and recommend that the two alternatives - connection oriented data transfer, and connectionless data transmission - be treated as complementary concepts, which can be applied in parallel to the different applications for which each is suited.

At the November, 1980 meeting of the ISO subcommittee responsible for OSI (TC97/SC16), a working party laid a solid foundation for this argument in two documents: Report of the Ad Hoc Group

on Connectionless Data Transmission[3], and Recommended Changes to Section 3 of [the Reference Model] to Include Connectionless Data Transmission[2]; and the importance of the issue was recognized by the full subcommittee in a resolution[25] calling for comments on the two documents from all member organizations. The question of how the connectionless data transmission concept should be reflected in the OSI architecture - and in particular, whether or not it should become an integral part of the Reference Model - will be debated again this summer, when the current Draft Proposed Standard Reference Model becomes a Draft International Standard. The remainder of this article will explore the issues that surround this question.

2 What Is Connectionless Data Transmission?

Connectionless data transmission (CDT), despite the unfamiliar name, is by no means a new concept. In one form or another, it has played an important role in the specification of services and protocols for over a decade. The terms "message mode"[ ], "datagram"[35], "transaction mode"[22,23,24], and "connection-free"[37,47] have been used in the literature to describe variations on the same basic theme: the transmission of a data unit in a single self-contained operation without establishing, maintaining, and terminating a connection.

Since connectionless data transmission and connection-oriented data transfer are complementary concepts, they are best understood in juxtaposition, particularly since CDT is most often defined by its relationship to the more familiar concept of a connection.

2.1 Connection-Oriented Data Transfer

A connection (or "(N)-connection", in the formal terminology of OSI) is an association established between two or more entities ("(N+1)-entities") for conveying data ("(N)-service-data-units"). The ability to establish (N)-connections, and to convey data units over them, is provided to (N+1)-entities by the (N)-layer as a set of services, called connection-oriented (N)-services. Connection-oriented interactions proceed through three distinct sequential phases: connection establishment; data transfer; and connection release. Figure 2 illustrates schematically the sequence of operations associated with connection-oriented interactions. In addition to this explicitly distinguishable duration, or "lifetime", a connection exhibits the following fundamental characteristics:

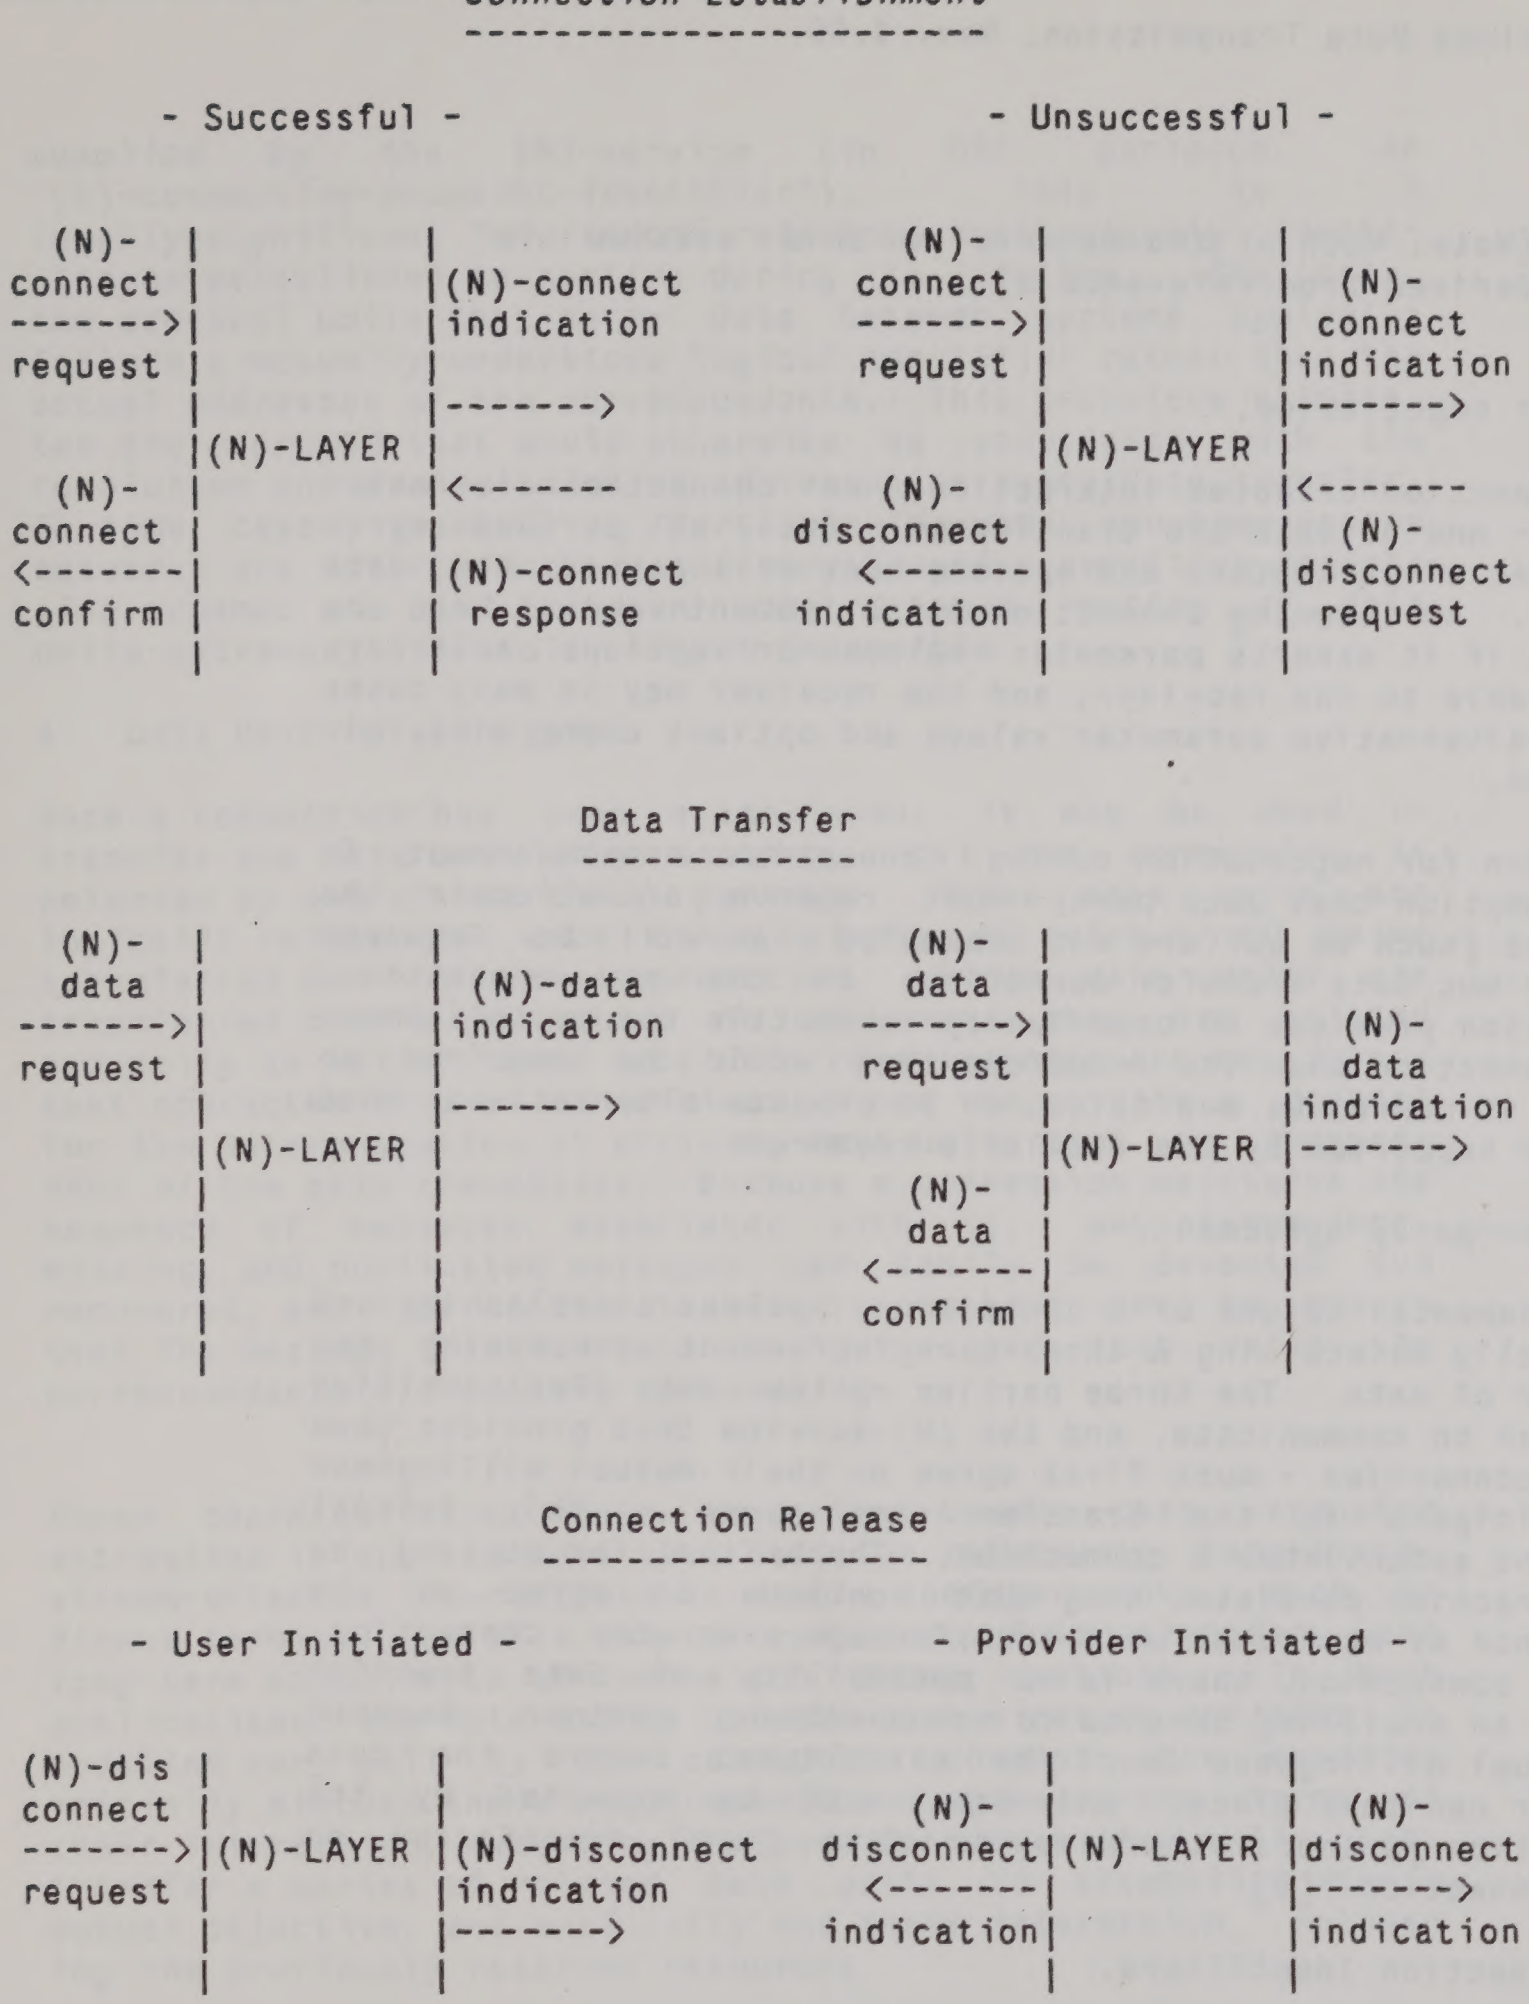

FIGURE 2 - Connection Oriented Interaction

[Note: Much of the material in this section is derived from reference 3]

1. Prior negotiation.

In a connection-oriented interaction, no connection is established - and no data are transferred - until all parties agree on the set of parameters and options that will govern the data transfer. An incoming connection establishment request can be rejected if it asserts parameter values or options that are unacceptable to the receiver, and the receiver may in many cases suggest alternative parameter values and options along with his rejection.

The reason for negotiation during connection establishment is the assumption that each party must reserve or allocate the resources (such as buffers and channels) that will be required to carry out data transfer operations on the new connection. Negotiation provides an opportunity to scuttle the establishment of a connection when the resources that would be required to support it cannot be dedicated, or to propose alternatives that could be supported by the available resources.

2. Three-party Agreement.

The fundamental nature of a connection involves establishing and dynamically maintaining a three-party agreement concerning the transfer of data. The three parties - the two (N+1)-entities that wish to communicate, and the (N)-service that provides them with a connection - must first agree on their mutual willingness to participate in the transfer (see above). This initial agreement establishes a connection. Thereafter, for as long as the connection persists, they must continue to agree on the acceptance of each data unit transferred over the connection. "With a connection, there is no possibility of data transfer through an unwilling service to an unwilling partner, because the mutual willingness must be established before the data transfer can take place, and data must be accepted by the destination partner; otherwise, no data [are] transferred on that connection."[3]

3. Connection Identifiers.

At connection establishment time, each participating (N+1)-entity is identified to the (N)-service by an (N)-address; the (N)-service uses these addresses to set up the requested connection. Subsequent requests to transfer data over the connection (or to release it) refer not to the (N)-address(es) of the intended recipient(s), but to a connection identifier

supplied by the (N)-service (in OSI parlance, an "(N)-connection-endpoint-identifier"). This is a locally-significant "shorthand" reference that uniquely identifies an established connection during its lifetime. Similarly, the protocol units that carry data between systems typically include a mutually-understood logical identifier rather than the actual addresses of the correspondents. This technique eliminates the overhead that would otherwise be associated with the resolution and transmission of addresses on every data transfer. In some cases, however - particularly when non-homogeneous networks are interconnected, and very location-sensitive addressing schemes are used - it can make dynamic routing of data units extremely difficult, if not impossible.

4. Data Unit Relationship.

Once a connection has been established, it may be used to transfer one data unit after another, until the connection is released by one of the three parties. These data units are logically related to each other simply by virtue of being transferred on the same connection. Since data units are transferred over a connection in sequence, they are related ordinally as well. These data unit relationships are an important characteristic of connections, since they create a context for the interpretation of arriving data units that is independent of the data themselves. Because a connection maintains the sequence of messages associated with it, out-of-sequence, missing, and duplicated messages can easily be detected and recovered, and flow control techniques can be invoked to ensure that the message transfer rate does not exceed that which the correspondents are capable of handling.

These characteristics make connection-based data transfer attractive in applications that call for relatively long-lived, stream-oriented interactions in stable configurations, such as direct terminal use of a remote computer, file transfer, and long-term attachments of remote job entry stations. In such applications, the interaction between communicating entities is modelled very well by the connection concept: the entities initially discuss their requirements and agree to the terms of their interaction, reserving whatever resources they will need; transfer a series of related data units to accomplish their mutual objective; and explicitly end their interaction, releasing the previously reserved resources.

2.2 Connectionless Data Transmission

In many other applications, however, the interaction between

entities is more naturally modelled by the connectionless data transmission concept, which involves the transmission of a single self-contained data unit from one entity to another without prior negotiation or agreement, and without the assurance of delivery normally associated with connection-based transfers. The users of a connectionless (N)-service may, of course, use their (N+1)-protocol to make any prior or dynamic arrangements they wish concerning their interpretation of the data transmitted and received; the (N)-service itself, however, attaches no significance to individual data units, and does not attempt to relate them in any way. Two (N+1)-entities communicating by means of a connectionless (N)-service could, for example, apply whatever techniques they might consider appropriate in the execution of their own protocol (timers, retransmission, positive or negative acknowledgements, sequence numbers, etc.) to achieve the level of error detection and/or recovery they desired. Users of a connectionless, as opposed to connection-oriented, (N)-service are not restricted or inhibited in the performance of their (N+1)-protocol; obviously, though, the assumption is that CDT will be used in situations that either do not require the characteristics of a connection, or actively benefit from the alternative characteristics of connectionless transmission.

Figure 3 illustrates schematically the single operation whereby a connectionless service may be employed to transmit a single data unit. Figure 4 shows a widely-implemented variation, sometimes called "reliable datagram" service, in which the service provider undertakes to confirm the delivery or non-delivery of each data unit. It must be emphasized that this is not a true connectionless service, but is in some sense a hybrid, combining the delivery assurance of connection-oriented service with the single-operation interface event of connectionless service.

Many of those involved in OSI standardization activities have agreed on a pair of definitions for connectionless data transmission, one for architectural and conceptual purposes, and one for service-definition purposes[4]. The architectural definition, which has been proposed for inclusion in the Reference Model, is:

"Connectionless Data Transmission is the transmission (not transfer) of an (N)-service-data-unit from a source (N)-service-access-point to one or more destination (N)-service-access-points without establishing an (N)-connection for the transmission."

The service definition, which is intended to provide a workable basis for incorporating a connectionless service into the

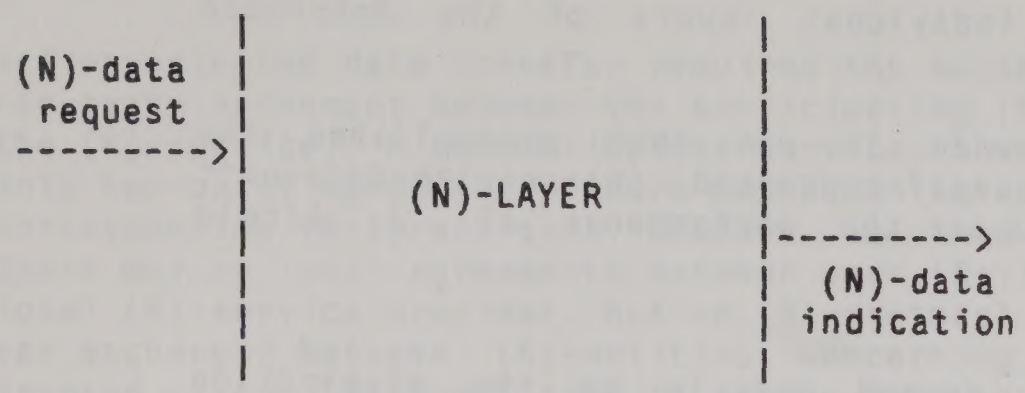

FIGURE 3 - Connectionless Data Transmission

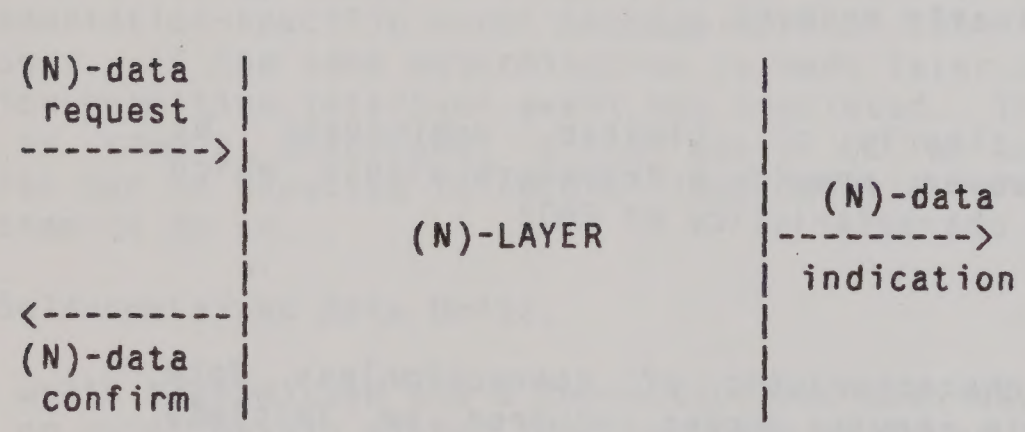

FIGURE 4 - "Reliable Datagram" Service

service descriptions for individual layers of the Reference Model, is:

"A Connectionless (N)-Service is one that accomplishes the transmission of a single self-contained (N)-service-data-unit between (N+1)-entities upon the performance of a single (N)-service access."

Both of these definitions depend heavily on the distinction between the terms "transmit", "transfer", and "exchange":

Transmit: "to cause to pass or be conveyed through space or a medium." This term refers to the act of conveying only, without implying anything about reception.

Transfer: "to convey from one place, person, or thing, to another." A one-way peer-to-peer connotation restricts the use of this term to cases in which the receiving peer is party to and accepts the data transferred.

Exchange: "to give and receive, or lose and take, reciprocally, as things of the same kind." A two-way peer-to-peer connotation restricts the use of this term to cases in which both give and receive directions are clearly evident.

These definitions are clearly of limited usefulness by themselves. They do, however, provide a framework within which to explore the following characteristics of CDT:

1. "One-shot" Operation.

The most user-visible characteristic of connectionless data transmission is the single service access required to initiate the transmission of a data unit. All of the information required to deliver the data unit - destination address, quality of service selection, options, etc. - is presented to the connectionless (N)-service provider, along with the data, in a single logical service-access operation that is not considered by the (N)-service to be related in any way to other access operations, prior or subsequent (note, however, that since OSI is not concerned with implementation details, the specific interface mechanism employed by a particular implementation of connectionless service might involve more than one interface exchange to accomplish what is, from a logical standpoint, a single operation). Once the service provider has accepted a data unit for connectionless transmission, no further communication occurs between the provider and the user of the service concerning the fate or disposition of the data.

2. Two-party Agreement.

Connection-oriented data transfer requires the establishment of a three-party agreement between the participating (N+1)-entities and the (N)-service. A connectionless service, however, involves only two-party agreements: there may be an agreement between the corresponding (N+1)-entities, unknown to the (N)-service, and there may be local agreements between each (N+1)-entity and its local (N)-service provider, but no (N)-protocol information is ever exchanged between (N)-entities concerning the mutual willingness of the (N+1)-entities to engage in a connectionless transmission or to accept a particular data unit.

In practice, some sort of a priori agreement (usually a system engineering design decision) is assumed to exist between the (N+1)-entities and the (N)-service concerning those parameters, formats, and options that affect all three parties as a unit. However, considerable freedom of choice is preserved by allowing the user of a connectionless service to specify most parameter values and options - such as transfer rate, acceptable error rate, etc. - at the time the service is invoked. In a given implementation, if the local (N)-service provider determines immediately (from information available to it locally) that the requested operation cannot be performed under the conditions specified, it may abort the service primitive, returning an implementation-specific error message across the interface to the user. If the same determination is made later on, after the service-primitive interface event has completed, the transmission is simply abandoned, since users of a connectionless service can be expected to recover lost data if it is important for them to do so.

3. Self-contained Data Units.

Data units transmitted via a connectionless service, since they bear no relationship either to other data units or to a "higher abstraction" (such as a connection), are entirely self-contained. All of the addressing and other information needed by the service provider to deliver a data unit to its destination must be included in each transmission. On the one hand, this represents a greater overhead than is incurred during the data transfer phase of a connection-oriented interaction; on the other, it greatly simplifies routing, since each data unit carries a complete destination address and can be routed without reference to connection-related information that may not, for example, be readily available at intermediate nodes.

4. Data Unit Independence.

The connectionless transmission of data creates no relationship, express or implied, between data units. Each invocation of a

connectionless service begins the transmission of a single data unit. Nothing about the service invocation, the transmission of the data by the connectionless service, or the data unit itself affects or is affected by any other past, present, or future operation, whether connection-oriented or connectionless. A series of data units handed one after the other to a connectionless service for delivery to the same destination will not necessarily be delivered to the destination in the same order; and the connectionless service will make no attempt to report or recover instances of non-delivery.

Note: A number of popular variations on CDT include features that run counter to those described above. These variations deserve to be discussed on their own merits, but should not be confused with the architectural concept of connectionless data transmission.

These characteristics make CDT attractive in applications that involve short-term request-response interactions, exhibit a high level of redundancy, must be flexibly reconfigurable, or derive no benefit from guaranteed in-sequence delivery of data.

3 The Rationale for Connectionless Data Transmission

Because CDT is not as widely understood as connection-oriented data transfer, it has often been difficult in the course of developing service and protocol definitions to adduce a rationale for incorporating CDT, and even more difficult to determine appropriate locations for connectionless service within the layered hierarchy of OSI. This section addresses the first concern; the next section will deal with the second.

The most natural way to discover the power and utility of the CDT concept is to examine applications and implementation technologies that depend on it. The following observations are distilled from the specifications and descriptions of actual protocols and systems (many of which have been implemented), and from the work of individuals and organizations engaged in the OSI standardization effort (quoted material is from reference 3, except where otherwise noted). They are divided into seven (occassionally overlapping) categories which classify the applications for which CDT is well suited.

Inward data collection involves the periodic active or passive sampling of a large number of data sources. A sensor net

gathering data from dedicated measurement stations; a network status monitor constantly refreshing its knowledge of a network environment; and an automatic alarm or security system in which each component regularly self-tests and reports the result, are all engaged in this type of interaction, in which a "large number of sources may be reporting periodically and asynchronously to a single reporting point. In a realtime monitoring situation, these readings could normally be lost on occassion without causing distress, because the next update would be arriving shortly. Only if more than one successive update failed to arrive within a specified time limit would an alarm be warranted. If, say, a fast connect/disconnect three-way handshake cost twice as much as a one-way [connectionless] data transmission which had been system engineered to achieve a certain acceptable statistical reliability figure, the cost of connection-oriented inward data collection for a large distributed application could be substantially greater than for [connectionless collection], without a correspondingly greater benefit to the user."[3]

Outward data dissemination is in a sense the inverse of the first category; it concerns the distribution of a single data unit to a large number of destinations. This situation is found, for example, when a node joins a network, or a commonly-accessible server changes its location, and a new address is sent to other nodes on the network; when a synchronizing message such as a real-time clock value must be sent to all participants in some distributed activity; and when an operator broadcasts a nonspecific message (e.g., "Network coming down in five minutes"). In such cases, the distribution cost (including time) may far exceed the cost of generating the data; controlling the overall cost depends on keeping the cost of dissemination as low as possible.

Request-response applications are those in which a service is provided by a commonly accessible server process to a large number of distributed request sources. The typical interaction consists of a single request followed by a single response, and usually only the highest-level acknowledgement - the response itself - is either necessary or meaningful. Many commercial applications (point of sale terminals, credit checking, reservation systems, inventory control, and automated banking systems) and some types of industrial process control, as well as more general information retrieval systems (such as videotex), fall into this category. In each case, the knowledge and expectation of each application component as to the nature of the interaction is represented in an application-process design and implementation that is known in advance, outside of OSI; lower level negotiations, acknowledgements, and other connection-related functions are often unnecessary and cumbersome.

An example of an application that combines the characteristics of inward data collection, outward data dissemination, and request-response interaction is described by the Working Group on Power System Control Centers of the IEEE Power Engineering Society in a recent letter to the chairman of ANSI committee X3T51 concerning the use of data communication in utility control centers[17]. They note that "a utility control center receives information from remote terminal units (located at substations and generating plants) and from other control centers, performs a variety of monitoring and control functions, and transmits commands to the remote terminals and coordinating information to other control centers." During the course of these operations, the following conditions occur:

1) Some measurements are transmitted or requested from remote terminals or control centers every few seconds. No attempt is necessarily made to recover data lost due to transmission error; the application programs include provisions for proper operation when input data is occassionally missing. [Inward data collection]

2) Some data items are transferred from commonly accessed remote sites or multi-utility pool coordination centers on a request-response basis. [Request-response interaction]

3) In some cases, an application program may require that some measurements be made simultaneously in a large number of locations. In these cases, the control center will broadcast a command to make th affected measurements. [Outward data dissemination]

In closing, they note that "utility control centers around the world use data communications in ways similar to those in the United States."

Broadcast and multicast (group addressed) communication using connection-oriented services is awkward at best and impossible at worst, notwithstanding the occassional mention of "multi-endpoint connections" in the Reference Model. Some characteristics of connection-based data transfer, such as sequencing and error recovery, are very difficult to provide in a broadcast/multicast environment, and may not even be desirable; and it is not at all easy to formulate a useful definition of broadcast/multicast acknowledgement that can be supported by a low-level protocol. Where group addressing is an important application consideration, connectionless data transmission is usually the only choice.

Certain special applications, such as digitized voice, data

telemetry, and remote command and control, involving a high level of data redundancy and/or real-time transmission requirements, may profit from the fact that CDT makes no effort to detect or recover lost or corrupted data. If the time span during which an individual datum is meaningful is relatively short, since it is quickly superceded by the next - or if, as in digitized voice transmission, the loss or corruption of one or even several data units is insignificant - the application might suffer far more from the delay that would be introduced as a connection-oriented service dealt with a lost or out-of-sequence data unit (even if retransmission or other recovery procedures were not invoked) than it would from the unreported loss of a few data units in the course of a connectionless exchange. Other special considerations - such as the undesirability, for security reasons, of maintaining connection-state information between data transfers in a military command and control system - add force to the argument that CDT should be available as an alternative to connection-oriented data transfer.

Local area networks (LANs) are probably the most fertile ground for connectionless services, which find useful application at several layers. LANs employ intrinsically reliable physical transmission media and techniques (baseband and broadband coaxial cable, fiber optics, etc.) in a restricted range (generally no greater than 1 or 2 kilometers), and are typically able to achieve extremely low bit error rates. In addition, the media-access contention mechanisms favored by LAN designers handle transmission errors as a matter of course. The usual approach to physical interconnection ties all nodes together on a common medium, creating an inherently broadcast environment in which every transmission can be received by every station. Taking advantage of these characteristics virtually demands a connectionless data link service, and in fact most current and proposed LANs - the Xerox Ethernet[43], the proposed IEEE 802 LAN standard[14,46], and many others - depend on such a service. As a bonus, because connectionless services are simpler to implement - requiring only two or three service primitives - inexpensive VLSI implementations are often possible.

In addition, the applications for which LANs are often installed tend to be precisely those best handled by CDT. Consider this list of eight application classes identified by the IEEE 802 Interface Subcommittee as targets for the 802 LAN standard[46]:

1. Periodic status reporting - telemetry data from instrumentation, monitoring devices associated with static or dynamic physical environments;

2. Special event reporting - fire alarms, overload or stressing conditions;

3. Security control - security door opening and closing, system recovery or initialization, access control;

4. File transfer;

5. Interactive transactions - reservation systems, electronic messaging and conferencing;

6. Interactive information exchange - communicating text and word processors, electronic mail, remote job entry;

7. Office information exchange - store and forward of digitized voice messages, digitized graphic/image handling;

8. Real-time stimulus and response - universal product code checkout readers, distributed point of sale cash registers, military command and control, and other closed-loop and real-time applications.

Of these, almost all have already been identified as classic examples of applications that have an essentially connectionless nature. Consider this more detailed example of (8): a local area network with a large number of nodes and a large number of services (e.g., file management, printing, plotting, job execution, etc.) provided at various nodes. In such a configuration, it is impractical to maintain a table at each node giving the address of every service, since changing the location of a single service would require updating the address table at every node. An alternative is to maintain a single independent "server lookup" service, which performs the function of mapping the name of a given service to the address of a server providing that service. The server-lookup server receives requests such as, "where is service X?", and returns the address at which an instance of service X is currently located. Communication with the server-lookup server is inherently self-contained, consisting of a single request/response exchange. Only the highest-level acknowledgement - the response from the lookup service giving the requested address - is at all significant. The native reliability of the local area network ensures a low error rate; if a message should be lost, no harm is done, since the request will simply be re-sent if a timely response does not arrive. Such an interaction is poorly modelled by the connection-oriented paradigm of opening a connection, transferring a stream of data, and closing the connection. It is perfectly suited to connectionless transmission techniques.

Network interconnection (internetworking) can be facilitated - especially when networks of different types are involved, as is often the case - if the internetwork service is connectionless;

and a number of related activities, such as gateway-to-gateway communication, exhibit the request-response, inward data collection, and outward data dissemination characteristics that are well supported by CDT. One of the best examples of a connectionless internetwork service is described in a document published by the National Bureau of Standards (Features of Internetwork Protocol[29], which includes a straightforward discussion of the merits of the connectionless approach:

> "The greatest advantage of connectionless service at the internet level is simplicity, particularly in the gateways. Simplicity is manifested in terms of smaller and less complicated computer code and smaller computer storage requirements. The gateways and hosts are not required to maintain state information, nor interpret call request and call clear commands. Each data-unit can be treated independently...Connectionless service assumes a minim[al] service from the underlying subnetworks. This is advantageous if the networks are diverse. Existing internet protocols which are intended for interconnection of a diverse variety of networks are based on a connectionless service [for example the PUP Internetwork protocol[44], the Department of Defence Standard Internet Protocol[31], and the Delta-t protocol developed at Lawrence Livermore Laboratory[45]]."

The principle motivating the development of internetwork services and protocols that make few assumptions about the nature of the individual network services (the "lowest common denominator" approach) was formulated by Carl Sunshine as the "local net independence principle"[39]: "Each local net shall retain its individual address space, routing algorithms, packet formats, protocols, traffic controls, fees, and other network characteristics to the greatest extent possible." The simplicity and robustness of connectionless internetworking systems guarantee their widespread use as the number of different network types - X.25 networks, LANs, packet radio networks, other broadcast networks, and satellite networks - increases and the pressures to interconnect them grow.

4 CDT and the OSI Reference Model

As a concept, connectionless data transmission complements the concept of connection-oriented data transfer throughout the OSI

architecture. As a basis for deriving standard OSI services and protocols, however, it has a greater impact on some layers of the Reference Model than on others. Careful analysis of the relative merits of connectionless and connection-oriented operation at each layer is necessary to control the proliferation of incompatible or useless options and preserve a balance between the power of the complementary concepts and the stabilizing objective of the OSI standardization effort.

Figure 5 illustrates the layered OSI hierarchy as it is most commonly represented (it shows two instances of the hierarchy, representing the relationship between two OSI systems). The following sections discuss the CDT concept in the context of each of the seven layers.

4.1 Physical Layer

The duality of connections and connectionless service is difficult to demonstrate satisfactorily at the physical layer, largely because the concept of a physical "connection" is both intuitive and colloquial. The physical layer is responsible for generating and interpreting signals represented for the purpose of transmission by some form of physical encoding (be it electrical, optical, acoustic, etc.), and a physical connection, in the most general sense (and restricting our consideration, as does the Reference Model itself, to telecommunications media), is a signal pathway through a medium or a combination of media. Is a packet radio broadcast network, then, using a "connectionless" physical service? No explicit signal pathway through a medium or media is established before data are transmitted. On the other hand, it can easily be argued that a physical connection is established with the introduction of two antennae into the "ether"; and if the antennae are aimed at each other and designed to handle microwave transmission, the impression that a physical connection exists is strengthened. Whether or not one recognizes the possibility of connectionless physical services - other than purely whimsical ones - will probably continue to depend on one's point of view, and will have no effect on the development of actual telecommunication systems.

4.2 Data Link Layer

Many data link technologies - particularly those coming into popular use with the growth of local area networking - are far easier to understand and work with when the traditional connection-oriented concepts (embodied, for example, in the widely-used HDLC, SDLC, and ADCCP standards) are replaced by the

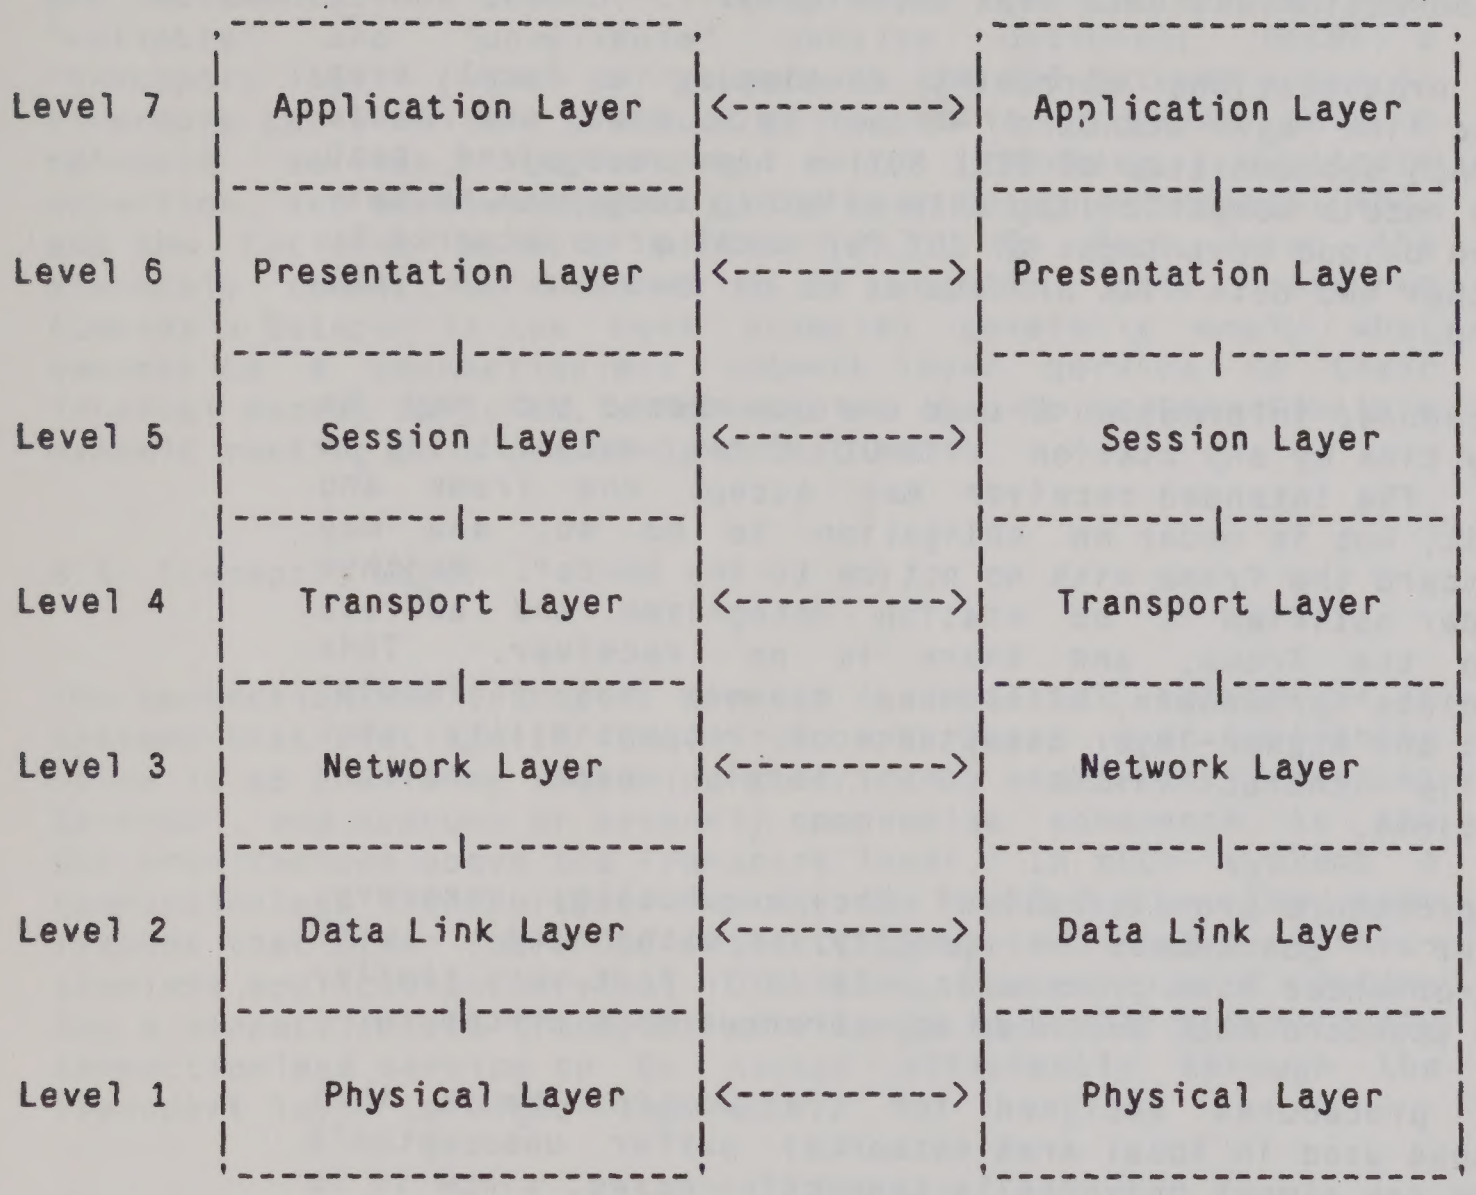

FIGURE 5 - Layered Hierarchy of Open Systems Interconnection

concept of connectionless data transmission. The previous discussion of local area networking has already made the point that the high-speed, short-range, intrinsically reliable broadcast transmission media used to interconnect stations in local area networks are complemented both functionally and conceptually by connectionless data link techniques.

One of the organizations currently developing a local area network data link layer standard - the Data Link and Media Access (DLMAC) subcommittee of IEEE 802 - has recognized both the need to retain compatibility with existing long-haul techniques and the unique advantages of CDT for local area networks by proposing that two data link procedures be defined for the IEEE 802 standard.

In one procedure, information frames are unnumbered and may be sent at any time by any station without first establishing a connection. The intended receiver may accept the frame and interpret it, but is under no obligation to do so, and may instead discard the frame with no notice to the sender. Neither is the sender notified if no station recognizes the address coded into the frame, and there is no receiver. This "connectionless" procedure, of course, assumes the "friendly" environment and higher-layer acceptance of responsibility that are usually characteristic of local area network implementations.

The other procedure provides all of the sequencing, recovery, and other guarantees normally associated with connection-oriented link procedures. It is in fact very similar to the ISO standard HDLC balanced asynchronous mode procedure.

Data link procedures designed for transmission media that (unlike those used in local area networks) suffer unacceptable error rates are almost universally connection-based, since it is generally more efficient to recover the point-to-point bit-stream errors detectable by connection-oriented data link procedures at the data link layer (with its comparatively short timeout intervals) than at a higher layer.

4.3 Network Layer

Connectionless network service is useful for many of the same reasons that were identified in the previous discussion of network interconnection: it greatly simplifies the design and implementation of systems; makes few assumptions about underlying services; and is more efficient than a connection-oriented service when higher layers perform whatever sequencing, flow control, and error recovery is required by user applications (in

fact, internetwork services are provided by the Network Layer). CDT also facilitates dynamic routing in packet- and message-switched networks, since each data unit (packet or message) can be directed along the most appropriate "next hop" unencumbered by connection-mandated node configurations. Examples of more or less connectionless network layer designs and implementations abound: Zilog's Z-net (which offers both "reliable" and "unreliable" service options); DECNET's "transport layer" (which corresponds to the OSI Network layer); Livermore Lab's Delta-t protocol (although it provides only a reliable service, performing error checking, duplicate detection, and acknowledgement); the User Datagram protocol[48]; and the Cyclades network protocol[38]. In fact, even the staunchly connection-oriented X.25 public data networks (Canada's Datapac is the best example) generally emply what amounts to a connectionless network-layer service in their internal packet switches, which enables them to perform flexible dynamic routing on a packet-by-packet basis.

4.4 Transport Layer

The connectionless transport service is important primarily in systems that distinguish the Transport layer and everything below it as providing something generically named the "Transport Service", and abandon or severely compromise adherence to the OSI architecture above the Transport layer. In such systems a connectionless transport service may be needed for the same reasons that other (more OSI-respecting) systems need a connectionless application service. Otherwise, the purpose of defining a connectionless transport service is to enable a uniformly connectionless service to be passed efficiently through the Transport layer to higher layers.

4.5 Session Layer

The whole notion of a session which binds presentation-entities into a relationship of some temporal duration is inherently connection-oriented. The purpose of defining a connectionless session service, therefore, is to enable a uniformly connectionless service to be passed efficiently through the session layer to higher layers. In this sense, the connectionless session service stands in precisely the same relationship to the connectionless transport service as a session-connection stands to a transport-connection.

4.6 Presentation Layer

Very much the same considerations apply to the Presentation layer as apply to the Session layer.

4.7 Application Layer

The most obvious reason to define a connectionless application service - to give user application processes access to the connectionless services of the architecture - is not the only one. The application layer performs functions that help user application processes to converse regarding the meaning of the information they exchange, and is also responsible for dealing with the overall system management aspects of the OSI operation. Over and above the many user-application requirements for connectionless service, it may be profitably employed by system management functions that monitor and report on the status of resources in the local open system; by application layer management functions that need to interact in a request-response mode with similar functions in other systems to perform security access control; and by user application process functions that monitor the status of activities in progress.

The potential availability of two complementary services at each layer of the architecture raises an obvious question - how to choose between them? It should be clear at this point that unilateral exclusion of one or the other, although it may simplify the situation for some applications, is not a general solution to the problem. There are actually two parts to the question: how to select an appropriate set of cooperative services for all seven layers during the design of a particular open system; and, if one or more layers of the system will offer both connection-oriented and connectionless services, how to provide for the dynamic selection of one or the other in a given circumstance.

The second part is easiest to dispose of, since actual systems - as opposed to the more abstract set of services and protocols collected under the banner of OSI - will generally be constructed in such a way as to combine services cooperatively, with some attention paid to the way in which they will interact to meet specific goals. Although two services may be provided at a given layer, logical combinations of services for different applications will generally be assembled according to relatively simple rules established during the design of the system.

Evaluating the requirements of the applications a system must

support and the characteristics of the preferred implementation technologies will also answer the first question. A system designed primarily to transport large files over a long-haul network would probably use only connection-oriented services. One designed to collect data from widely scattered sensors for processing at a central site might provide a connectionless application service but use a connection-oriented network service to achieve compatibility with a public data network. Another system, built around a local area network bus or ring, might use a connectionless data link service regardless of the applications supported; if several LANs sere to be interconnected, perhaps with other network types, it might also employ a connectionless internetwork service.

The definition of OSI standard services and protocols, however, must consider the general case, so as to accomodate a wide range of actual-system configurations. The motivating principle should be to achieve a balance between the two goals of power and simplicity. The service definition for each layer must include both connection-oriented and connectionless services; otherwise, the utility of a service at one layer could be negated by the unavailability of a corresponding service elsewhere in the hierarchy. However, the role played by each service may be radically different from one layer to the next. The Presentation, Session, and Transport layers, for instance, need to support their respective connectionless services only because the Application layer, which must provide a connectionless service to user applications, cannot do so effectively if they do not. Recognizing these role variations opens up the possibility of restoring a measure of the simplicity lost in the introduction of choice at each layer by limiting, not the choices, but the places in the hierarchy where conversion from one choice to the other - connection to connectionless, or vice versa - is allowed (see figure 6). At this stage in the development of the CDT concept, it appears that there are exscellent reasons for allowing such a conversion to take place in the Application, Transport, and Network layers (and in the Data Link layer, if some physical interconnection strategies are deemed to be connectionless). In the other layers, the provision of one kind of service to the next-higher layer must always be accomplished by using the same kind of service from the next-lower layer (see figure 7). (This principle of like-to-like mapping is not related to multiplexing; it refers to service types (connection-oriented and connectionless), not to actual services.) Adopting such a restriction would contribute to the achievement of the balance mentioned above, without excluding those combinations of services that have demonstrated their usefulness.

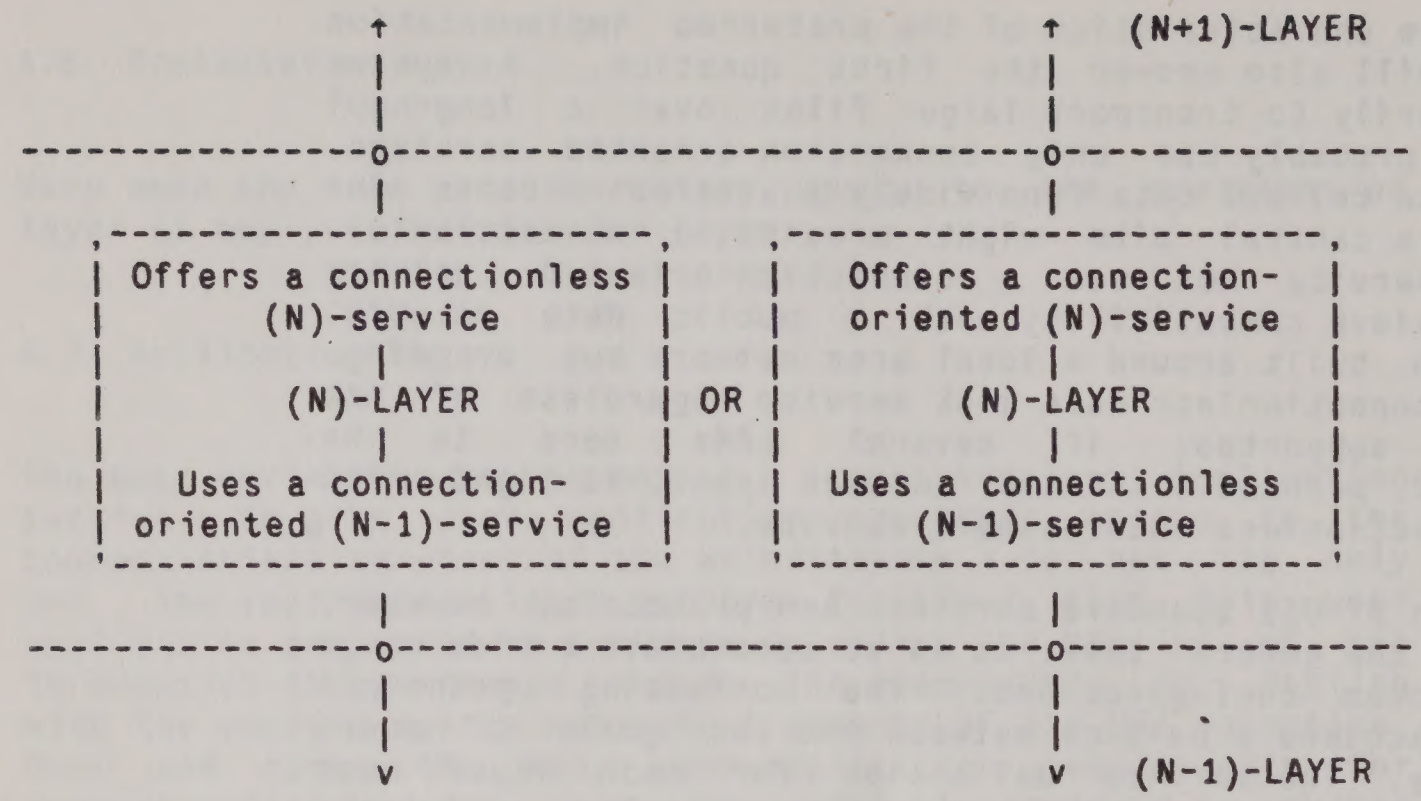

FIGURE 6 - Service Type Conversion

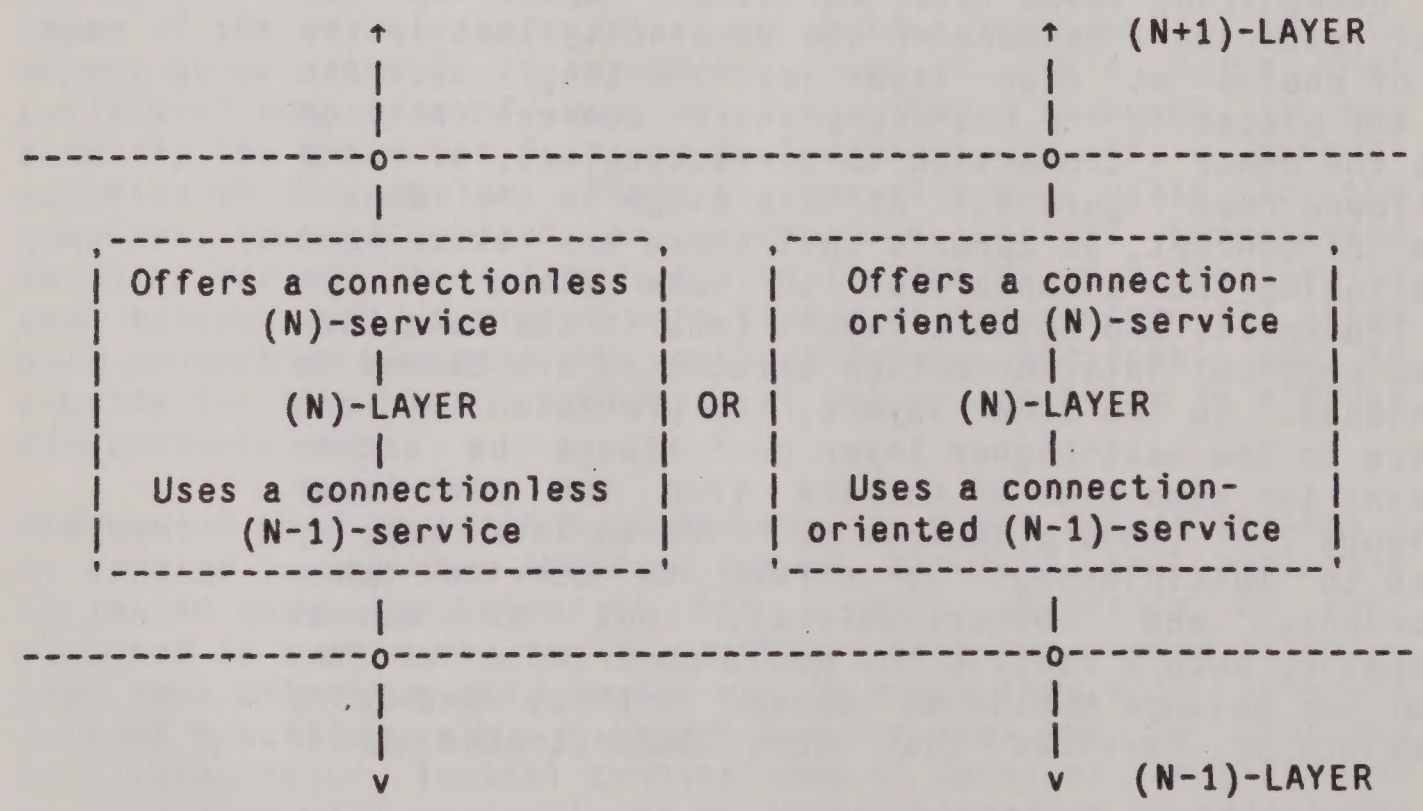

FIGURE 7 - Same-Service Mapping

5 Summary

Support for incorporating connectionless data transmission as a basic architectural element of the Reference Model has grown as understanding of the concept has become more widespread. The protocol development sponsored by various agencies of the U.S. Department of Defense, for example, have long recognized connections and connectionless transmission as complementary concepts, and have employed both. Similar work being carried out by a division of the Institute for Computer Science and Technology at the National Bureau of Standards, the result of which will be a series of Federal Information Processing Standards, depends heavily on connectionless as well as connection-oriented concepts. The importance of CDT to some of these U.S. efforts is reflected in comments received by ANSI committee X3T5 during the recent Reference Model ballot period, one of which states that "Publication of this material [DP7498] without incorporation of the concerns associated with Connectionless Data Trans[mission] makes a mockery of U.S. interests."[18] A somewhat less emotional expression of the same sentiment is embodied in the official U.S. Position on Connectionless Data Transmission[9], in which X3T5, the responsible U.S. organization, "endorses SC16/N555 [Recommended Changes to Section 3 of [the Reference Model] to Include CDT] without exception and announces its intention to pursue vigorously the incorporation of CDT as the first major extension to the Basic Reference Model of OSI." In the same document, X3T5 notes that it "intends to issue and maintain a version of DP7498 to be referred to as DP7498-prime, incorporating the CDT extensions." That there is also significant international support for the CDT concept is clear, however, from the membership of the ISO SC16/WG1 Ad Hoc Group on Connectionless Data Transmission, which produced the N555 document last November; it includes representatives from France, Japan, Germany, and the United Kingdom as well as from the U.S. Those who believe that the CDT concept is an essential part of the OSI architecture hope that eventually the DP7498-prime document, or its successor, will replace the exclusively connection-oriented Reference Model before the latter becomes an International Standard.

6 Acknowledgements

[to be supplied]

APPENDIX A - Vocabulary

OSI Terminology

The following terms are defined in either the text or the vocabulary annex (or both) of the Draft Proposed Reference Model of OSI (ISO/DP7498). Some terms are given more than one definition in different sections of the Reference Model; these are marked with an asterisk (*), to indicate that selection of the accompanying definition involved the author's personal judgement.

[to be supplied]

(N)-connection
(N)-service-access-point
(N)-service-access-point-address
(N)-layer
system
(N)-entity
(N)-connection-endpoint-identifier

CDT Terminology

The following terms, not yet part of the standard OSI vocabulary, relate to the concept of connectionless data transmission.

"Connectionless Data Transmission is the transmission (not transfer) of an (N)-service-data-unit from a source (N)-service-access-point to one or more destination (N)-service-access-points without establishing an (N)-connection for the transmission."

"A Connectionless (N)-Service is one that accomplishes the

transmission of a single self-contained (N)-service-data-unit between (N+1)-entities upon the performance of a single (N)-service access."

Transmit: "to cause to pass or be conveyed through space or a medium." This term refers to the act of conveying only, without implying anything about reception.

Transfer: "to convey from one place, person, or thing, to another." A one-way peer-to-peer connotation restricts the use of this term to cases in which the receiving peer is party to and accepts the data transferred.

Exchange: "to give and receive, or lose and take, reciprocally, as things of the same kind." A two-way peer-to-peer connotation restricts the use of this term to cases in which both give and receive directions are clearly evident.

datagram
unit-data transfer/transmission
transaction (from SC1/N688)
data transmission (from DIS 2382 Section 9)

[End of Appendix A]

APPENDIX B - References

1. Data Processing - Open Systems Interconnection - Basic Reference Model.

Source: ISO/TC97/SC16
Reference: ISO/DP7498
X3T51/80-67
X3S33/X3T56/80-121
X3S37/80-115
Date: 12/80

2. Recommended Changes to Section 3 of 97/16 N537, Basic Specifications of the Reference Model of OSI, to Include Connectionless Data Transmission.

Source: ISO/TC97/SC16/WG1 Ad Hoc Group on Connectionless Data Transmission
Reference: ISO/TC97/SC16/N555
X3S37/81-9
X3T51/80-68
X3S33/X3T56/80-122
Date: 11/80

3. Report of the Ad Hoc Group on Connectionless Data Transmission.

Source: ISO/TC97/SC16/WG1 Ad Hoc Group on Connectionless Data Transmission
Reference: ISO/TC97/SC16/N566
X3T51/80-69
X3S33/X3T56/81-13
X3S37/81-35
Date: 11/80

4. Definitions of the Term "Connectionless Data Transmission" (a letter to the chairman of ANSC X3T51 from the acting chairman of ANSC X3T56).

Source: ANSC X3S33/X3T56
Reference: X3S33/X3T56/81-22
X3T51/81-2
X3S37/81-6
Date: 1/81

5. Connectionless Provisions for OSI Reference Model.

Source: ANSC X3S37
Reference: ISO/TC97/SC6/WG2/W12
X3S37/81-16R
Date: 2/81

6. Comments on Recommended Changes to Section 3 of 97/16 N537, Basic Specification of the Reference Model of OSI, to include Connectionless Data Transmission, SC16/N555.

Source: DIN (FRG)
Reference: ISO/TC97/SC6/WG2/W10
Date: 2/81

7. Connectionless Data Transmission.

Source: X3S33/X3T56 Ad Hoc Group on Connectionless Data Transmission
Reference: X3S33/X3T56/81-26
Date: 1/81

8. Contribution to Document ISO/TC97/SC16 N555 Concerning the Extension of General Concepts from the Basic Reference Model to Connectionless Data Transfer Mode.

Source: ISO/TC97/SC16/WG1 Ad Hoc Model Extension Group B
Reference:
Date: 3/81

9. US Position on Connectionless Data Transmission.

Source: ANSC X3T5
Reference: ISO/TC97/SC16/N605
X3T51/81-26
Date: 3/81

10. Revision of SC16/N551 to Include Connectionless Data Transmission.

Source: ANSC X3S33/X3T56
Reference: ISO/TC97/SC16/N602
X3S33/X3T56/81-67
X3T51/81-20
X3S37/81-17
Date: 3/81

11. Report of USA Vote and Comments on ISO DP7498.

Source: ANSC X3T5
Reference: ISO/TC97/SC16/N590
X3T51/81-29
Date: 3/81

12. USA Proposed Revision to Draft Basic Session Service Specification, ISO TC97/SC16 N553.

Source: ANSC X3S33/X3T56
Reference: ISO/TC97/SC16/N597
X3S33/X3T56/81-39R
X3T51/81-28
Date: 3/81

13. USA Proposed Revision to Draft Transport Service Specification, ISO TC97/SC16 N563.

Source: ANSC X3S33/X3T56
Reference: ISO/TC97/SC16/N601
X3S33/X3T56/81-33R
X3T51/81-17
Date: 3/81

14. Comments on Connectionless Data Transmission.

Source:	Robert F. Stover, Honeywell Inc.
Reference:	Private communication
Date:	4/81

15. Proposed Changes to the OSI Transport Layer.

Source:	Gregory Ennis, Sytek Inc.
Reference:	X3T51 Reference Model Editing Group V3.B
Date:	3/81

16. Review of the ISO Draft Proposal (DP 7498), Open System Interconnection Reference Model (Project IPSC-0168).

Source:	National Security Agency, Central Security Service, Department of Defense
Reference:	NSA/CSS Serial T095/008/81
	X3T51 Reference Model Editing Group V3.F
Date:	3/81

17. Comments on Draft Proposal ISO/DP7498.

Source:	Working Group on Power System Control Centers, IEEE Power Engineering Society
Reference:	X3T51 Reference Model Editing Group V3.I, V4.4
Date:	3/81

18. Review of ISO Draft Proposal 7498 (Open Systems Interconnection).

Source:	Department of the Air Force
Reference:	X3T51 Reference Model Editing Group V3.J, V4.5, V1.15, V2.H
Date:	3/81

19. Proposed Improvements to Section 6 of DP7498.

Source: A. Lyman Chapin, Data General Corporation
Reference: X3T51 Reference Model Editing Group V3.M
Date: 3/81

20. Comments on Section 7.4 of DP7498.

Source: ANSC X3S33/X3T56
Reference: X3S33/X3T56/81-30
X3T51 Reference Model Editing Group V3.H
Date: 3/81

21. Comments on DP7498.

Source: ANSC X3S33/X3T56
Reference: X3S33/X3T56/81-60
X3T51 Reference Model Editing Group V3.N
Date: 3/81

22. USA Position Concerning Progression of the Reference Model of Open Systems Interconnection (Parts I and II of USA Comments on N309).

Source: ANSC X3T5
Reference: ISO/TC97/SC16/N405
X3T5/80-120
X3T51/80-43
Date: 9/80

23. Addenda to the USA Position Concerning Progression of OSI Reference Model (Parts I and II).

Source: ANSC X3T5
Reference: X3T5/80-143
X3T51/80-63
Date: 9/80

24. US Position on the WG1 Rapporteur's Report of October 1980.

Source: ANSC X3T5
Reference: X3T5/80-142
X3T51/80-62
Date: 10/80

25. Resolutions: ISO/TC97/SC16 - Open Systems Interconnection: Berlin - November 12 - 14, 1980.

Source: ISO/TC97/SC16
Reference: ISO/TC97/SC16/N570
X3S33/X3T56/80-11
Date: 11/80

26. NBS Analysis of Major US Government Requirements of Transport Protocol Services.

Source: National Bureau of Standards, US Department of Commerce
Reference: ISO/TC97/SC16/N404
X3T51/80-32
X3S33/X3T56/80-82
Date: 9/80

27. Features of the Transport and Session Protocols.

Source: National Bureau of Standards, US Department of Commerce
Reference: X3S33/X3T56/80-30
Date: 3/80

28. Specification of the Transport Protocol.

Source: National Bureau of Standards, US Department of Commerce
Reference: X3S33/X3T56/81-59
Date: 2/81

29. Features of Internetwork Protocol.

Source: National Bureau of Standards, US Department of Commerce
Reference: X3T51/81-23
X3S33/X3T56/80-96
X3S37/81-31
Date: 7/80

30. Service Specification of an Internetwork Protocol.

Source: National Bureau of Standards, US Department of Commerce
Reference: X3T51/81-24
X3S33/X3T56/81-18
X3S37/81-32
Date: 9/80

31. DoD Standard Internet Protocol.

Source: US Department of Defense Advanced Research Projects Agency
Reference: X3S33/X3T56/80-17
X3S37/80-17
Date: 1/80

32. Connectionless Data Transfer (letter from the chairman of X3T51 to X3T55, X3T56, and X3S3).

Source: John Day, Digital Technology, Inc.
Reference: X3T51/80-76
Date: 12/80

33. Local Area Networks and the OSI Reference Model.

Source: Robert R. Shatzer, Hewlett-Packard Corp.
Reference: X3T51/80-38
Date: 8/80

34. An Introduction to Local Area Networks.

Source: David D. Clark, et. al.
Reference: IEEE Proceedings 66:11
Date: 11/78

35. Issues in Packet-Network Interconnection.

Source: V.G. Cerf and P.T. Kirstein
Reference: IEEE Proceedings 66:11
Date: 11/78

36. Connectionless Data Transfer.

Source: John Neumann, Microdata Corp.
Reference: X3S33/X3T56/80-120
Date: 12/80

37. A Protocol for Packet Network Interconnection.

Source: V.G. Cerf and R.E. Kahn
Reference: IEEE Transactions on Communication COM-22 No. 5
Date: 5/74

38. The CYCLADES End-to-End Protocol.

Source: H. Zimmermann
Reference: Proceedings of the IEEE Vol. 66 No. 11
Date: 11/78

39. Interprocess Communication Protocols for Computer Networks.

Source: Carl Sunshine, USC/ISI
Reference: Stanford Digital Systems Laboratory TR105
Date: 12/75

40. CCITT Recommendation X.25 - Interface Between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) for Terminals Operating in the Packet Mode on Public Data Networks.

Source: CCITT Study Group VII
Reference: COM VII/489
Date: 11/80

41. An Analysis of ARPAnet Protocols.

Source:
Reference:
Date:

42. ISO High-Level Data Link Control - Elements of Procedure.

Source: ISO
Reference: ISO/IS4335
Date: 1977

43. ETHERNET Specification (Version 1.0)

Source: Xerox Corporation
Reference: X3T51/80-50
Date: 9/80

44. PUP: An Internetwork Architecture.

Source: D.R. Boggs, J.F. Shoch, E.A. Taft, R.M. Metcalfe
Reference: IEEE Transactions on Communications COM-28 No. 4
Date: 4/80

45. Delta-t Protocol Preliminary Specification.

Source: R.W. Watson
Reference: Lawrence Livermore Laboratories
Date: 11/79

46. The Evolving IEEE 802 (Local Network) Standard.

Source: Bryan R. Hoover, Hewlett-Packard Corporation
Reference:
Date:

47. A System for Interprocess Communication in a Resource Sharing Computer Network.

Source: D. Walden
Reference: Communications of the ACM Vol. 15
Date: 4/72